arXiv:1407.7508v1 [cs.LG] 28 Jul 2014

Efficient Regularized Regression for Variable Selection with L_0 Penalty

Zhenqiu Liu[1], Gang Li[2]

[1]Samuel Oschin Comprehensive Cancer Institute, Cedars-Sinai Medical Center, Los Angeles, CA 90048, USA. [2]Department of Biostatistics, School of Public Health, University of California at Los Angeles, Los Angeles, CA 90095-1772, U.S.A.

Abstract

Variable (feature, gene, model, which we use interchangeably) selections for regression with high-dimensional BIGDATA have found many applications in bioinformatics, computational biology, image processing, and engineering. One appealing approach is the L_0 regularized regression which penalizes the number of nonzero features in the model directly. L_0 is known as the most essential sparsity measure and has nice theoretical properties, while the popular L_1 regularization is only a best convex relaxation of L_0. Therefore, it is natural to expect that L_0 regularized regression performs better than LASSO. However, it is well-known that L_0 optimization is NP-hard and computationally challenging. Instead of solving the L_0 problems directly, most publications so far have tried to solve an approximation problem that closely resembles L_0 regularization.

In this paper, we propose an efficient EM algorithm (L_0EM) that directly solves the L_0 optimization problem. L_0EM is efficient with high dimensional data. It also provides a natural solution to all L_p $p \in [0, 2]$ problems, including LASSO with $p = 1$, elastic net with $p \in [1, 2]$, and the combination of L_1 and L_0 with $p \in (0, 1]$. The regularized parameter λ can be either determined through cross-validation or AIC and BIC. Theoretical properties of the L_0-regularized estimator are given under mild conditions that permit the number of variables to be much larger than the sample size. We demonstrate our methods through simulation and high-dimensional genomic data. The results indicate that L_0 has better performance than LASSO and L_0 with AIC or BIC has similar performance as computationally intensive cross-validation. The proposed algorithms are efficient in identifying the non-zero variables with less-bias and selecting biologically important genes and pathways with high dimensional BIGDATA.

1 Introduction

Variable selection with regularized regression has been one of the hot topics in machine learning and statistics. Regularized regressions identify outcome associated features and estimate nonzero parameters simultaneously, and are particularly useful for high-dimensional BIGDATA with small sample sizes. In many real applications, such as bioinformatics, image and signaling processing, and engineering, a large number of features are measured, but only a small number of features are associated with the dependent variables. Including irrelevant variables in the model will lead to overfitting and deteriorate the prediction performance. Therefore, different regularized regression methods have been proposed for variable selection and model construction. L_0 regularized regressions, which directly penalize the number of non-zero parameters, are the most essential sparsity measure. Several popular information criteria, including Akaike information criterion (AIC) (Akaike 1974), Bayesian information criterion (BIC) (Schwarz 1978), and risk inflation criteria (RIC) (Foster and George 1994), are based on L_0 penalty and have been used extensively for variable selections. However, solving a general L_0 regularized optimization is NP hard and computational challenging. Exhaustive search with AIC or BIC over all possible combinations of features is computationally infeasible with high-dimensional BIGDATA.

Different alternatives have been proposed for the regularized regression problem. One common approach is to replace L_0 by L_1. L_1 is known as the best convex relaxation of L_0. L_1 regularized regression (Tibshirani 1996) is convex and can be solved by an efficient gradient decent algorithm. Minimizing L_1 is equivalent to minimizing L_0 under certain conditions. However, the estimates of L_1 regularized regression are asymptotically biased, and LASSO may not always choose the true model consistently (Zou 2006). Experimental results by Mancera and Portilla (2006) also posed additional doubt about the equivalence of minimizing L_1 and L_0. Moreover, there were theoretical results (Lin et al., 2010) showing that while L_1 regularized regression never outperforms L_0 by a constant, in some cases L_1 regularized regression performs infinitely worse than L_0. Lin et al.(2010) also showed that the optimal L_1 solutions are often inferior to L_0 solutions found using greedy classic stepwise regression, although solutions with L_1 penalty can be found effectively. More recent approaches aimed

to reduce bias and overcome discontinuity include SCAD (Fan and Li, 2001), L_p $p \in (0, 1]$ regularization (Liu et al., 2010; Mazumder et al., 2011), and MC+ (Zhang, 2010). Even though there are some effects for solving the L_0 regularized optimization problems (Dicker et al., 2012; Lu & Zhang, 2013), L_0 was either approximated by a continuous smooth function, or transformed into a much larger ranking optimization problem. To the best of our knowledge, there is no method that optimizes L_0 directly.

In this paper, we propose an efficient EM algorithm (L_0EM) that directly solves the L_0 regularized regression problem. L_0EM effectively deals with L_0 optimization by solving a sequence of convex L_2 optimizations and is efficient for high dimensional data. It also provides a natural solution to all L_p $p \in [0, 2]$ problems, including LASSO with $p = 1$, elastic net with $p \in [1, 2]$ (Zou & Zhang 2009), and the combination of L_1 and L_0 with $p \in (0, 1]$ (Liu & Wu, 2007). While the regularized parameter λ for LASSO must be tuned through cross-validation, which is time-consuming, the optimal λ with L_0 regularized regression can be pre-determined with different model selection criteria such as AIC, BIC and RIC. We demonstrate our methods through simulation and high-dimensional genomic data. The proposed methods identify the non-zero variables with less-bias and outperform the LASSO method by a large margin. They can also choose the biologically important genes and pathways effectively.

2 Methods

Given a $n \times 1$ dependent variable $\mathbf{y}$, and an $n \times m$ feature matrix X, a linear model is defined as

$$\mathbf{y} = X\theta + \varepsilon,$$

where n is the number of samples and m is the number of variables and $n \ll m$, $\theta = [\theta_1, \ldots, \theta_m]^t$ are the m parameters to be estimated, and $\varepsilon \sim N(0, \sigma^2 I_n)$ are the random errors with mean 0 and variance σ^2. Assume only a small subset of $\{\mathbf{x}_j\}_{j=1}^m$ has nonzero θ_js. Let $R \subseteq \{1, \ldots, m\}$ be the subset index of relevant variables with $\theta_j \neq 0$, and $O \subseteq \{1, \ldots, m\}$ be the index of irrelevant features with 0 coefficients, we have $R \cup O = \{1, 2, \ldots, m\}$, $X_R \cup X_O =$

X, and $\theta_R \cup \theta_O = \theta$, where $\theta_O = 0$. The error function for L_1 regularized regression is

$$E = \frac{1}{2}||\mathbf{y} - X\theta||^2 + \frac{\lambda}{2}||\theta||_0 = \frac{1}{2}\sum_{i=1}^{n}(y_i - \sum_{j=1}^{m}\theta_j x_{ij})^2 + \frac{\lambda}{2}\sum_{j=1}^{m} I(\theta_j \neq 0), \tag{1}$$

where $||\theta||_0 = \sum_{j=1}^{m} I(\theta_j \neq 0) = |R|$ counts the number of nonzero parameters. One observation is that equation (1) is equivalent to the following equation (2), when reaching the optimal solution.

$$E = \frac{1}{2}||\mathbf{y} - X\theta||^2 + \frac{\lambda}{2}||\theta||_0 = \frac{1}{2}||\mathbf{y} - X\theta||^2 + \frac{\lambda}{2}\sum_{j\in R} 1 = \frac{1}{2}||\mathbf{y} - X\theta||^2 + \frac{\lambda}{2}|R|, \tag{2}$$

because θ_O is a zero vector. Our L_0EM methods will be derived from equation (2). We can rewrite equation (2) as the following two equations:

$$E = \frac{1}{2}||\mathbf{y} - X\theta||^2 + \frac{\lambda}{2}\sum_{j\in R}\frac{\theta_j^2}{\eta_j^2} \tag{3}$$

$$\eta = \theta. \tag{4}$$

Given η_j, equation (3) is a convex quadratic function and can be optimized by taking the first order derivative:

$$\frac{\partial E}{\partial \theta_R} = \lambda\theta_R \oslash \eta_R^2 - X_R^t(\mathbf{y} - X\theta) = 0, \tag{5}$$

$$\text{and} \quad \frac{\partial E}{\partial \theta_O} = \mathbf{0}, \quad \text{as} \quad \theta_O = \eta_O = \mathbf{0}, \tag{6}$$

where $\oslash$ indicates element-wise division. Rewriting (5) and (6), we have

$$\lambda\theta_R - \eta_R^2 \odot X_R^t(\mathbf{y} - X\theta) = 0, \tag{7}$$

$$\lambda\theta_O - \eta_O^2 \odot X_O^t(\mathbf{y} - X\theta) = 0, \ \ \forall \ \ \lambda > 0 \tag{8}$$

where $\odot$ is element-wise multiplication, $\eta_R^2 \odot X_R^t = [\eta_R^2 \odot \mathbf{x}_{1R}^t, \ldots, \eta_R^2 \odot \mathbf{x}_{nR}^t]$, and $\eta_O^2 \odot X_O^t = [\eta_O^2 \odot \mathbf{x}_{1O}^t, \ldots, \eta_O^2 \odot \mathbf{x}_{nO}^t] = \mathbf{0}$. Let $X_\eta^t = \eta^2 \odot X^t$ and combining equations (7) and (8) together, we have

$$\eta^2 \odot \frac{\partial E}{\partial \theta} = \lambda\theta - \eta^2 \odot X^t(\mathbf{y} - X\theta) = \lambda\theta - X_\eta^t(\mathbf{y} - X\theta) = 0. \tag{9}$$

Solving Equation (9), we have the following explicit solution.

$$\theta = (X_\eta^t X + \lambda I)^{-1} X_\eta^t \mathbf{y} \tag{10}$$

$$\eta = \theta, \tag{11}$$

where equation (10) can be considered as the M-step of the EM algorithm maximizing $-E$, and equation (11) can be regarded as the E-step with $E(\eta) = \theta$. Equations (10) and (11) together can also be treated as a fixed point iteration method in nonlinear optimization.

Theorem 2.1. *Given an input matrix X, output matrix $\mathbf{y}$, and initialized solution θ^0, the nonlinear system determined by equations (10) and (11) will converge to a unique solution, as long as the regularized parameter $0 < \lambda < ||X_\theta^t X||$, and the estimated solution is closer to the true solution after each iterative EM step.*

Proof: Equations (10) and (11) are the same as:

$$\theta = (X_\theta^t X + \lambda I)^{-1} X_\theta^t \mathbf{y} = (\theta^2 \odot X^t X + \lambda I)^{-1} (\theta^2 \odot X^t) \mathbf{y} \tag{12}$$

First, $G(\theta) = (\theta^2 \odot X^t X + \lambda I)^{-1} (\theta^2 \odot X^t) \mathbf{y}$ is Lipschitz continuous for $\theta \in R^m$, and

$$\begin{aligned} \frac{\partial G(\theta)}{\partial \theta} &= (\theta^2 \odot X^t X + \lambda I)^{-2} [(\theta^2 \odot X^t X + \lambda I)(2\theta \odot X^t)\mathbf{y} - 2\theta \odot X^t X (\theta^2 \odot X^t)\mathbf{y}] \\ &= (\theta^2 \odot X^t X + \lambda I)^{-2} [2\lambda\theta \odot X^t \mathbf{y}] = 2\lambda (\theta^2 \odot X^t X + \lambda I)^{-1} \mathbf{1}_m \\ &= 2\lambda (X_\theta^t X + \lambda I)^{-1} \mathbf{1}_m, \end{aligned} \tag{13}$$

where I is the identity matrix and $\mathbf{1}_m = [1, \ldots, 1]^t$ is a m-dimensional vector of 1s, and we substitute equation (12) into equation (13) to get the result.

Because $\lambda < ||X_\theta^t X||$, it is clear from equation (13) that

$$\left| \frac{\partial G(\theta)}{\partial \theta_j} \right| = \frac{2\lambda}{||X_\theta^t X|| + \lambda} < \frac{2\lambda}{\lambda + \lambda} = 1,$$

$\forall\ j = 1, \ldots, m$. Therefore, there is a Lipschitz constant

$$\gamma = \left|\left| \frac{\partial G(\theta)}{\partial \theta} \right|\right|_\infty = \max_j \left\{ \left| \frac{\partial G(\theta)}{\partial \theta_j} \right| \right\} < 1.$$

Now given the initial value for equations (10) and (11) $\eta = \theta^0 \in R^m$, the sequence $\{\theta^r\}$ remains bounded because $\forall\ i = 1, \ldots, r$,

$$||\theta^{i+1} - \theta^i||_\infty = ||G(\theta^i) - G(\theta^{i-1})||_\infty \simeq ||\frac{\partial G(\theta)}{\partial \theta}(\theta^i - \theta^{i-1})||_\infty \leq \left|\left|\frac{\partial G(\theta)}{\partial \theta}\right|\right|_\infty ||\theta^i - \theta^{i-1}||_\infty$$
$$= \gamma||\theta^i - \theta^{i-1}||_\infty \leq \ldots \leq \gamma^i||\theta^1 - \theta^0||_\infty.$$

and therefore

$$||\theta^r - \theta^0||_\infty = ||\sum_{i=0}^{r-1}(\theta^{i+1} - \theta^i)||_\infty \leq ||\theta^1 - \theta^0||_\infty \sum_{i=0}^{r-1}\gamma^i$$
$$\leq \frac{||\theta^1 - \theta^0||_\infty}{(1-\gamma)}.$$

Now $\forall\ r, k \geq 0$,

$$||\theta^{r+k} - \theta^r||_\infty = ||G(\theta^{r+k-1}) - G(\theta^{r-1}||_\infty \leq \gamma||\theta^{r+k-1} - \theta^{r-1}||_\infty$$
$$\leq \gamma||G(\theta^{r+k-2}) - G(\theta^{r-2})||_\infty \leq \gamma^2||\theta^{r+k-2} - \theta^{r-2}||_\infty$$
$$\leq \ldots \leq \gamma^r||\theta^k - \theta^0||_\infty \leq \frac{\gamma^r||\theta^1 - \theta^0||_\infty}{1-\gamma}.$$

Hence,

$$\lim_{r,k\to\infty} ||\theta^{r+k} - \theta^r||_\infty = 0,$$

and therefore $\{\theta^r\}$ is a Cauchy sequence that has a limit solution θ^*.

Next the uniqueness of the solution is easy to show. Assuming there were two solutions θ^* and $\theta^\diamond$, then

$$||\theta^* - \theta^\diamond||_\infty = ||G(\theta^*) - G(\theta^\diamond)||_\infty \leq \gamma|\theta^* - \theta^\diamond||_\infty. \tag{14}$$

Since $\gamma < 1$, equation (14) can only hold, if $||\theta^* - \theta^\diamond||_\infty = 0$. i.e. $\theta^* = \theta^\diamond$, so the solution of the EM algorithm is unique.

Finally, the EM algorithm will be closer to the true solution at each step, because

$$||\theta^{r+1} - \theta^*||_\infty = ||G(\theta^r) - G(\theta^*||_\infty \leq \gamma||\theta^r - \theta^*||_\infty.$$

Lemma 2.1. *Assuming that relevant features are independent, i. e.* $\mathbf{x}_i^t\mathbf{x}_j = 0$, $\forall\ i \neq j$ & $i, j \in \{1, \ldots, m\}$, *then the maximal regularized parameter* λ *can be determined by*

$$\lambda_{\max} = \max\left\{\frac{(\mathbf{x}_j^t\mathbf{y})^2}{4\mathbf{x}_j^t\mathbf{x}_j}\right\}_{j=1}^m$$

Proof: For each feature $\mathbf{x}_j$ and corresponding coefficient θ_j, equations (9) and (11) can be rewritten as

$$\eta_j^2 \frac{\partial E(\theta)}{\partial \theta_j} = \lambda\theta_j - \eta_j^2 \mathbf{x}_j^t(\mathbf{y} - X\theta) = 0$$
$$\eta_j = \theta_j, \quad \forall \ j \in \{1, \ldots, m\}.$$

The above two equations are the same as:

$$\lambda\theta_j - \theta_j^2 \mathbf{x}_j^t(\mathbf{y} - X\theta) = 0. \tag{15}$$

If $\theta_j = 0$, then any $\lambda > 0$ will satisfy equation (15). On the other hand, if $\theta_j \neq 0$, because $\mathbf{x}_i^t \mathbf{x}_j = 0$, equation (15) becomes the following quadratic equation:

$$(\mathbf{x}_j^t \mathbf{x}_j)\theta_j^2 - (\mathbf{x}_j^t \mathbf{y})\theta_j + \lambda = 0. \tag{16}$$

One necessary condition for equation (16) to have a solution is:

$$(\mathbf{x}_j^t \mathbf{y})^2 - 4(\mathbf{x}_j^t \mathbf{x})\lambda \geq 0, \quad \Rightarrow \quad \lambda \leq \frac{(\mathbf{x}_j^t \mathbf{y})^2}{4(\mathbf{x}_j^t \mathbf{x}_j)}$$

Therefore the maximal λ is

$$\lambda_{\max} = \max \left\{ \frac{(\mathbf{x}_j^t \mathbf{y})^2}{4\mathbf{x}_j^t \mathbf{x}_j} \right\}_{j=1}^{m}.$$

If $\lambda > \lambda_{\max}$, equation (15) holds only if all $\theta_j = 0$, $\forall\ j = 1, \ldots, m$.

Both Theorem 2.1 and Lemma 1 2.1 provide some useful guidance for implementing the method and choosing the regularized parameter λ. Theorem 2.1 shows that the EM algorithm always converges to a unique solution, given a certain λ and initial solution θ^0, and the estimated value is closer to the true solution after each EM iteration. Note that different initial values may still reach different solution, because of the non-convex L_0 penalty. Therefore, it is critical to choose a good initial value. Our experiences with the method indicate that initializing with the estimates from L_2 based ridge regression will usually lead to quick converge and super performance. The EM algorithm is as follows.

The L_0EM Algorithm:

Given a $0 < \lambda \leq \lambda_{\max}$, small numbers ϵ and ε,

and training data $\{X, \mathbf{y}\}$,

Initializing $\theta = (X^tX + \lambda I)^{-1}X^t\mathbf{y}$,

While 1,

E-step: $\eta = \theta$

M-step: $X_\eta^t = \eta^2 \odot X^t = [\eta^2 \odot \mathbf{x}_1, \ldots, \eta^2 \odot \mathbf{x}_n^t]$

$\theta = (X_\eta^t X + \lambda I)^{-1}X_\eta^t\mathbf{y}$

if $||\theta - \eta|| < \varepsilon$, Break; End

End

$\theta(|\theta| < \epsilon) = 0$.

Similar procedures can be extended to general L_p, $p \in [0, 2]$ without much difficulty. L_p based EM algorithm L_pEM is reported in Appendix.

Consistency and Oracle Property: Let θ_0 be the true parameter value. The following conditions will be used later for theoretical properties of the L_0-regularized estimator of θ_0.

CONDITIONS

(C1) $ln(m) = o(n)$ as $n \to \infty$.

(C2) There exists a constant $K > 0$ such that $\lambda_{max}(\frac{X^tX}{n}) \leq K < \infty$ for large n, where for any matrix B, $\lambda_{max}(B)$ denotes the largest eigenvalue of B.

(C3) $\frac{\max_j ||\mathbf{x}_j||}{\sqrt{n}} = O(\sqrt{ln(mn)})$ or $O(1)$ as $n, m \to \infty$.

(C4) There exists a constant $c > 0$ such that $\frac{\min_j ||\mathbf{x}_j||^2}{n} \geq c > 0$ for large n, m.

(C5) $\mu(X) \equiv \max_{1 \leq i < j \leq m} \frac{|\mathbf{x}_i^t \mathbf{x}_j|}{||\mathbf{x}_i|| \cdot ||\mathbf{x}_j||} = O(\sqrt{\frac{ln(m)}{n}})$.

(C6) $||\theta_0||_0 = O(1)$.

The above conditions are very mild. Condition (C1) trivially holds for $m \leq n$ and for $m > n$. In particular, (C1) is satisfied even for ultra-high dimensional case such as $m = exp(n^\alpha)$ for $0 < \alpha < 1$. (C2) is a standard condition for linear regression. Chi (2013, Section 3.2) gives

examples satisfying(C3)-(C4). For example, (C3) and (C4) trivially hold if $||\mathbf{x}_i = \sqrt{n}$ for all $j = 1, \ldots, m$. (C5) is referred to as the coherence condition under which the covariates are not highly colinear; see Bunea et al. (2007), Candes and Plan (2009), and Chi (2013). (C6) implies that the model is sparse.

The following theorem is a direct consequence of Chi (2013).

Theorem 2.2 (*Consistency*). *Assume that conditions (C1)-(C6) hold. Let* $n(\nu) = (1 - \nu)[1 + 1/\mu(X)]$ *for some* $0 < \nu < 1$. *For any* $0 < q < \frac{1}{2}$, *let* $\lambda = \frac{3ln(m/q)}{\nu[1+\mu(X)]} \frac{\max_j ||\mathbf{x}_j||^2}{\min_j ||\mathbf{x}_j||^2}$ *and*

$$\hat{\theta} = arg \min_{||\theta||_0 \leq n(\nu)} E_n(\theta),$$

Then, with probability tending to 1,

$$||\hat{\theta} - \theta_0|| = O_p(\sqrt{\frac{\ln(nm)}{n}}) \tag{17}$$

Proof Note that the normal linear model in this paper is a special case of the exponential model of Chi (2009): $p_t(y) = \exp(ty - \Lambda(t))$ with $t = \frac{\mathbf{x}_i^t \theta}{\sigma^2}$ and $\Lambda(t) = \frac{\sigma^2 t^2}{2}$. Then, (17) follows immediately from Theorem 3.1 of Chi (2009).

Model Recovery: Next we show that L_0-regularized regression recovers the true model under mild conditions.

Theorem 2.3 (*Oracle Property*). *Assume that conditions (C1)-(C6) hold. Let* $A = \{1 \leq j \leq m : \theta_{0j} \neq 0\}$, *and* $A^c = \{1, 2, \ldots, m\} \backslash A$. *Then, the minimizer* $\hat{\theta}$ *in Theorem 2.2 must satisfy* $\hat{\theta}_j = 0$ *for* $j \in A^c$.

Proof Let $\alpha_n = \sqrt{\frac{\ln(nm)}{n}}$. For any θ such that $||\theta - \theta_0|| < C\alpha_n$ for some constant $C > 0$ and $\sum_{j \in A^c} I(\theta_j \neq 0) \geq 1$, let

$$\tilde{\theta}_j = \begin{cases} \theta_j & \text{if} \quad j \in A \\ 0 & \text{if} \quad j \in A^c \end{cases}$$

Then,

$$
\begin{aligned}
&E_n(\theta) - E_n(\tilde{\theta}) \\
&= \frac{1}{2n}(\theta-\tilde{\theta})^T X^T X(\theta-\tilde{\theta}) - \frac{1}{n}(\theta-\tilde{\theta})^T X^T(y - X\tilde{\theta}) + \frac{\lambda}{2}(||\theta||_0 - ||\tilde{\theta}||_0) \\
&= \frac{1}{2n}(\theta-\tilde{\theta})^T X^T X(\theta-\tilde{\theta}) - \frac{1}{n}(\theta-\tilde{\theta})^T X^T(X\theta_0 + \epsilon - X\tilde{\theta}) + \frac{\lambda}{2}(||\theta||_0 - ||\tilde{\theta}||_0) \\
&= \frac{1}{2}(\theta-\tilde{\theta})^T \left(\frac{X^T X}{n}\right)(\theta-\tilde{\theta}) - (\theta-\tilde{\theta})^T \left(\frac{X^T X}{n}\right)(\theta_0-\tilde{\theta}) + \\
&+ \frac{1}{\sqrt{n}}(\theta-\tilde{\theta})^T \cdot \frac{1}{\sqrt{n}} X^T\epsilon + \frac{\lambda}{2}(||\theta||_0 - ||\tilde{\theta}||_0) \\
&= I_1 + I_2 + I_3 + I_4
\end{aligned}
$$

Because $||\tilde{\theta} - \theta_0|| \leq ||\theta - \theta_0||$, we have $\theta - \tilde{\theta} = O(\alpha_n)$. Thus, $I_1 = O(\alpha_n^2)$ and $I_2 = O(\alpha_n^2)$. Moreover,

$$
\left|\left|\frac{1}{\sqrt{n}}\epsilon^t X\right|\right| = O_p(\sqrt{k\sigma^2}), \qquad \text{as } n \to \infty
$$

where $k = rank(X) \leq n$. Hence,

$$
|I_3| \leq \frac{1}{\sqrt{n}}||\theta - \tilde{\theta}|| \cdot \left|\left|\frac{1}{\sqrt{n}} X^T\epsilon\right|\right| = O(\alpha_n) \cdot O_p(\sqrt{k/n}) = O_p(\alpha_n).
$$

Furthermore,

$$
\begin{aligned}
I_4 &= \frac{\lambda}{2}(||\theta||_0 - ||\tilde{\theta}||_0) \\
&= \frac{\lambda}{2}\sum_{j=1}^{m}[I(\theta_j \neq 0) - I(\tilde{\theta}_j \neq 0)] \\
&= \frac{\lambda}{2}\left[\sum_{j\in A} 0\right] + \frac{\lambda}{2}\sum_{j\in A^c}[I(\theta_j \neq 0) - 0] \\
&= \frac{\lambda}{2}\sum_{j\in A^c} I(\theta_j \neq 0) \geq \frac{\lambda}{2} \cdot 1 > 0.
\end{aligned}
$$

By conditions (C3)-C(5), $\lambda = O(ln(m) \cdot ln(nm))$. Therefore, the first three terms I_1, I_2 and I_3 are dominated by λ in probability as $n \to \infty$. Therefore, with probability tending to 1,

$$
E_n(\theta) - E_n(\tilde{\theta}) > 0. \tag{18}
$$

This completes the proof of Theorem 2.3.

Determination of λ: The regularized λ determines the sparsity of the model. The standard approach for choosing λ is cross-validation and the optimal λ is determined by the minimal mean squared error (MSE) of the test data ($MSE = \sum(y_i - \hat{y_i})^2/n$). One could also adapt the stability selection (SS) approach for λ determination (Liu et al.2010; Meinshausen, 2010). It chooses the smallest λ that minimizes the inconsistences in number of nonzero parameters with cross-validation. We first calculate the mean and standard deviation (SD) of the number of nonzero parameters for each λ, and then find the smallest λ with 0 SD, where 0 SD indicates that all models in k-fold cross validation has the same number of nonzero estimates. Our experiences indicate that the larger λ chosen from both minimal MSE and stability selection ($\lambda = \max\{\lambda_{mse}, \lambda_{ss}\}$) has the best performance. Choosing optimal λ from cross-validation is computationally intensive and time consuming. Fortunately, unlike LASSO, identifying the optimal λ for L_0 does not require to use cross validation. The optimal λ_{opt} can be determined by variable selection criteria. The optimal λ_{opt} can be directly picked using AIC, BIC, or RIC criteria with $\lambda_{opt} = 2$, $\log n$, or $2\log m$, respectively. Each of these criteria is known to be optimal under certain conditions. This is a huge advantage of L_0, especially for BIGDATA problems.

3 Simulations

To evaluate the performance of L_0 and L_1 regulation, we assume a linear model $\mathbf{y} = X\theta + \varepsilon$, where the input matrix X is from Gaussian distribution with mean $\mu = \mathbf{0}$ and different covariance structures Σ, where $\Sigma(i,j) = r^{|i-j|}$ with $r = 0, 0.3, 0.6, 0.8$ respectively. The true model is $\mathbf{y} = 2\mathbf{x}_1 - 3\mathbf{x}_2 + 4\mathbf{x}_5 + \varepsilon$ with $\varepsilon \sim N(0,1)$. Therefore, only three features are associated with output $\mathbf{y}$, and the rest of the θ_is are zero. In our first simulation, we first compare L_0 and L_1 regularized regression with a relative small number of features $m = 50$ and a sample size of n =100. Five-fold cross validation is used to determine the optimal λ and compare the model performance. We seek to fit the regularized regression models over a range of regularization parameters λ. Each λ is chosen from $\lambda_{\min} = 1e-4$, to $\lambda_{\max}$ with 100 equally log-spaced intervals, where $\lambda_{\max} = \max\{X^t\mathbf{y}\}$ for L_1 and $\max\left\{\frac{(\mathbf{x}_j^t\mathbf{y})^2}{4\mathbf{x}_j^t\mathbf{x}_j}\right\}$ for L_0. Lasso function in the statistics toolbox of MATLAB (www.mathworks.com) is used

r	L_0			L_1		
	# SF	MSE	$\|\|\hat{\theta}-\theta\|\|$	# S.F.	MSE	$\|\|\hat{\theta}-\theta\|\|$
0	3.39($\pm$1.1)	1.01($\pm$0.14)	0.206($\pm$0.12)	14.5($\pm$3.45)	1.19($\pm$0.19)	0.38($\pm$0.1)
0.3	3.37($\pm$0.9)	1.02($\pm$0.16)	0.23($\pm$0.12)	14.5($\pm$2.91)	1.21($\pm$0.19)	0.41($\pm$0.19)
0.6	3.49($\pm$1.7)	1.02($\pm$0.23)	0.23($\pm$0.16)	13.5($\pm$3.0)	1.26($\pm$0.2)	0.54($\pm$0.15)
0.8	3.32($\pm$0.9)	1.06($\pm$0.15)	0.28($\pm$0.21)	11.7($\pm$2.69)	1.3($\pm$0.21)	0.89($\pm$0.25)

Table 1: Performance measures for L_0 and L_1 regularized regression over 100 simulations, where values in the parenthesis are the standard deviations, and # SF: number of average selected features; MSE: Average mean squared error; $||\hat{\theta}-\theta||$: average absolute bias when comparing true and estimated parameters.

for comparison. Cross-validation with MSE is implemented nicely in the toolbox. The computational results are reported in Table 1. Table 1 shows that L_0 outperforms LASSO in all categories by a substantial margin, when using the popular test MSE measure for model selection. In particular, the number of variables selected by L_0 are very close to the true number of variables (3), while LASSO selected more than 11 features on average with different correlation structures (r = 0, 0.3, 0.6, 0.8). The test MSEs and bias both increase with the growth of correlation among features for both L_0 and LASSO, but the test MSE and bias of L_0 are substantially lower than these of LASSO. The maximal MSE of L_0 is 1.06, while the smallest MSE of L_1 is 1.19, and the largest bias of L_0 is 0.28, while the smallest bias of LASSO is 0.38. In addition (results are not shown in Table 1), L_0 correctly identifies the true model 81, 74, 81, and 82 times for r = 0, 0.3, 0.6, and 0.8 respectively over 100 simulations, while LASSO never chooses the correct model. Therefore, compared to L_0 regularized regression, LASSO selects more features than necessary and has larger bias in parameter estimation. Even though it is possible to get a correct model with LASSO using a larger λ, the estimated parameters will have a bigger bias and worse predicted MSE.

The same parameter setting is used for our second simulation, but the regularized parameter λ is determined by the larger λ from both minimal MSE and stability selection ($\lambda = \max\{\lambda_{MSE}, \lambda_{SS}\}$). The computational results are reported in Table 2. Table 2 shows that the average number of associated features is much closer to 3 with sightly larger test

r	L_0			L_1		
	# SF	MSE	$\|\|\hat{\theta}-\theta\|\|$	# S.F.	MSE	$\|\|\hat{\theta}-\theta\|\|$
0	3.09($\pm$0.53)	1.04($\pm$0.15)	0.18($\pm$0.11)	13.3($\pm$4.56)	1.21($\pm$0.17)	0.39($\pm$0.1)
0.3	3.08($\pm$0.54)	1.04($\pm$0.15)	0.17($\pm$0.07)	14.5($\pm$4.20)	1.22 ± 0.17)	0.42($\pm$0.19)
0.6	3.10($\pm$0.46)	1.07($\pm$0.17)	0.21($\pm$0.10)	13.8($\pm$5.4)	1.27($\pm$0.47)	0.57($\pm$0.25)
0.8	3.02($\pm$0.14)	1.04($\pm$0.14)	0.26($\pm$0.13)	13.4($\pm$4.91)	1.25($\pm$0.21)	0.74($\pm$0.25)

Table 2: Performance measures for L_0 and L_1 regularized regression with $\lambda = \max\{\lambda_{mse}, \lambda_{ss}\}$ over 100 simulations, where values in the parenthesis are the standard deviations, and # SF: number of average selected features; MSE: Average mean squared error; $||\hat{\theta}-\theta||$: average absolute bias when comparing true and estimated parameters.

MSEs. The maximal average number of features is 3.1 with $r = 0.6$, reduced from 3.49 with the test MSE only. In fact, with this combined model selection criteria and 100 simulations, L_0EM identified the true model with three nonzero parameters 95, 95, 95, and 97 times respectively (not shown in the table), while LASSO did not choose any correct models. The average bias of the estimates with L_0EM is also reduced. These indicate that the combination of test MSE and stability selection in cross-validation leads to better model selection results than MSE alone with L_0EM. However, the computational results did not improve much with LASSO. Over 13 features on average were selected under different correlation structures, suggesting that LASSO inclines to select more spurious features than necessary. A much more conservative criteria with larger λ is required to select the right number of features, which will induce larger MSE and bias, and deteriorate the prediction performance.

Simulation with high- dimensional data

Our third simulation deals with high-dimensional data with the number of samples $n = 100$, and the number of features $m = 1000$. The correlation structure is set to $r = 0, 0.3, 0.6$, and the same model $\mathbf{y} = 2\mathbf{x}_1 - 3\mathbf{x}_2 + 4\mathbf{x}_5 + \varepsilon$ was used for evaluating the performance of L_0 and L_1. The simulation was repeated 20 times. The computational results are reported in Table 3. Table 3 shows that L_0 outperforms LASSO by a large margin when correlations among

	Measures	r = 0	r = 0.3	r = 0.6
L_0	# SF	3(±0)	2.9(±0.47)	2(±0.73)
	$\|\|\hat{\theta} - \theta\|\|$	0.14(±0.09)	0.39(±0.63)	1.69(±1.25)
	Test MSE	1.14(±0.34)	1.59(±1.3)	2.8(±1.72)
	# True Model	20/20	15/20	5/20
L_1	# SF	24(±18.4)	31.3(±20.7)	36.7(±16.5)
	$\|\|\hat{\theta} - \theta\|\|$	0.57(±0.11)	0.73(±0.13)	1.14(±0.25)
	Test MSE	1.50(±0.25)	1.63(±0.29)	1.92(±0.41)
	# True Model	0/20	0/20	0/20

Table 3: Performance measures for L_0 and L_1 regularized regression with cross validation and $\lambda = \max\{\lambda_{mse}, \lambda_{ss}\}$ over 20 simulations and the sample size of $n = 100$, and $m = 1000$, where values in the parenthesis are the standard deviations, and # SF: number of average selected features; MSE: Average mean squared error; $||\hat{\theta} - \theta||$: average absolute bias when comparing true and estimated parameters.

features are low. When there is no correlation among features, 20 out of 20 simulations identify the true model with L_0, and 15 out of 20 simulations choose the correct model when $r = 0.3$, while LASSO again chooses more features than necessary and no true model was found under any correlation setting. However, when correlations among features are large with $r = 0.6$, the results are mixed. L_0 can still identify 5 out of 20 correct models, but the test MSE and bias of the parameter estimate of L_0 are slightly large than those of LASSO. In addition, we notice that L_0 is a more sparse model when correlation increases, indicating that L_0 tends to choose independent features. The regularization path of L_0 regression is shown in Figure 1. As shown in the top panel of Figure 1, the three associated features first increase their values when λ goes larger, and then go to zero when λ becomes extremely big, while the rest of the irrelevant features all go to zero when λ increases. Unlike LASSO, which shrinks all parameters uniformly, L_0 will only forces the estimates of irrelevant features go to zero, while keep the estimates of relevant features to their true value. This is the well-known Oracle property of L_0. For this specific simulation, the three parameters $[\hat{\theta_1}, \hat{\theta_2}, \hat{\theta_5}] = [1.85, -2.94, 4.0]$, very close to their true values $[2, -3, 4]$. The middle and bottom panels are the test MSE and the standard deviation of the number of

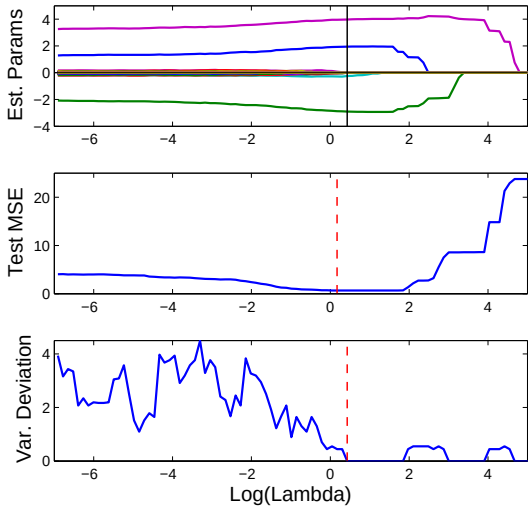

Figure 1: Regularized path for L_0 penalized regression with n=100, m =1000, and r = 0.3

nonzero variables. The optimal λ is chosen from the the larger λ with minimal test MSE and stability selection as shown in the vertical lines of Figure 1.

L_0 regularized regression without cross validation

Choosing the optimal parameter λ_{opt} with cross-validation is time consuming, especially with BIGDATA. As we mentioned previously, the optimal λ can be picked from theory instead of cross validation. Since we are dealing with the $n \ll m$ BIGDATA problem, RIC with $\lambda_{opt} = 2 \log m$ is penalized too much for such problem. So computational results with AIC and BIC without cross validation are reported in Table 4. Table 4 shows that L_0 regularized regression with AIC and BIC performs very well, when compared with the results from computationally intensive cross-validation in Table 3. Without correlation, BIC identifies the true model (100%), which is the same as cross-validation in Table 3, and better than AIC's 78%. The bias of BIC (0.16) is only slightly higher than that of cross-validation (0.14), but lower than that of AIC (0.19). Even though MSE*s with AIC and BIC are in-sample mean squared errors, which are not comparable to the test MSE with cross validation, larger MSE* with BIC indicates that BIC is an more stringent criteria than AIC and selects less variables. With mild correlation (r = 0.3) and some sacrifices in bias and MSE* , BIC seems to perform the best in variable selection, since the average number of features selected

	Measures	r = 0	r = 0.3	r = 0.6
AIC	# SF	3.26(±0.54)	3.72(±1.94)	4.8(±2.77)
	$\|\|\hat{\theta} - \theta\|\|$	0.19(±0.09)	0.36(±0.58)	1.02(±1.2)
	MSE*	0.96(±0.14)	1.02(±0.31)	1.27(±0.51)
	# True Model	78/100	73/100	59/100
BIC	# SF	3.0(±0.0)	3.0(±0.38)	2.89(±0.80)
	$\|\|\hat{\theta} - \theta\|\|$	0.16(±0.08)	0.45(±0.69)	1.80(±1.20)
	MSE*	0.97(±0.15)	1.29(±0.81)	2.48(±1.17)
	# True Model	100/100	94/100	53/100

Table 4: Performance measures for L_0 regularized regression with AIC and BIC over 100 simulations with $n = 100$, and $m = 1000$, where values in the parenthesis are the standard deviations, and # SF: number of average selected features; MSE*: In-sample average mean squared error; $||\hat{\theta} - \theta||$: average absolute bias when comparing true and estimated parameters.

is exactly 3 and 94% of the simulations recognize the true model, while AIC chooses more features (3.72) than necessary and only 73% of the simulations are right on targets. Cross validation is the most tight measure with 2.9 features on average and 75% of the simulations finding the correct model. When the correlations among the variables are high (r= 0.6), the results are mixed. Both BIC and AIC correctly identify more than half of the true models, while cross validation only recognizes 25% (5/20) of the model correctly. Therefore, comparing with the computationally intensive cross validation, both BIC and AIC perform reasonable well. The computational results of BIC is comparable to the results of cross validation, while the computational time is only 1/500 of the time for cross validation, if the free-parameter λ_{opt} is chosen from 100 candidate λs with 5-fold cross validation.

Simulations for graphical models

One important application of L_0 regularized regression is to detect high-order correlation structures, which has numerous real-world applications including gene network analysis. Given a matrix X, letting $\mathbf{x}_j$ be the jth variable, and X_{-j} be the remaining variables, we have

$P(\mathbf{x}_j|X_{-j}) \sim N(X_{-j}\theta, \sigma^2)$, where the coefficients θ measures the partial correlations between $\mathbf{x}_j$ and the rest variables. Therefore, the high-order structure of X has been determined via a series of L_1 regularized regression for each $\mathbf{x}_j$ with the remaining variables X_{-j} (Peng et al.2009; Liu & Ihler, 2011). The collected regression nonzero coefficients are the edges on the graph. The drawback of such approach is computationally intensive, because the regularized parameter λ for L_1 have to be determined through cross validation. For instance, given a matrix X with 100 variables, to find the optimal λ_{opt} from 100 candidate λs with 5-fold cross validation, 500 models need to be evaluated for each variable $\mathbf{x}_j$. Therefore a total of $500 \times 100 = 50000$ models have to be estimated to detect the dependencies among X with LASSO. It usually takes hours to solve this problem. However, only 100 models are required to identify the same correlation structure with L_0 regularized regression and AIC or BIC. Solving such a problem with L_0 without cross-validation only takes less than one minute. Finally, negative correlations between genes are difficult to confirm and seemingly less biologically relevant (Lee et al., 2004). Most national databases are constructed with similarity (dependency) measures. it is straight forward to study only the positive dependency by simply setting $\theta(\theta < 0) = 0$ in the EM algorithm.

We simulate two network structures similar to those in Zhang & Mallick (2013) (i) Band 1 network, where Σ is a covariance matrix with $\sigma_{ij} = 0.6^{|i-j|}$, so $A = \Sigma^{-1}$ has a band 1 network structure, and (ii) A more difficult problem for a Band 2 network with weaker correlations, where $A = -\Sigma^{-1}$ with $a_{ij} = \begin{cases} 0.25 & \text{if } |i-j| = 1, \\ 0.4 & \text{if } |i-j| = 2, \\ 0 & \text{Otherwise.} \end{cases}$ The sample sizes are n = 50, 100, and 200, respectively and the number of variables is $m = 100$. L_0 regularized regression with AIC and BIC is used to detect the network (correlation) structure. The consistence between the true and predicted structures is measured by the area under the ROC curve (AUC), false discovery (positive) rate (FDR/FPR), and false negative rate (FNR) of edges. The computational results are shown in Table 5. Table 5 shows that both AIC and BIC performed well. Both achieved at least 0.90 AUC for Band 1 network and 0.8 AUC for Band 2 network with different sample sizes. AIC performed slightly better than BIC, especially for Band 2 network with weak correlations and small sample sizes. This is reasonable because

	AIC			BIC		
Band 1	AUC	FDR(%)	FNR (%)	AUC	FDR (%)	FNR (%)
$n = 50$	.95(±.01)	.29(±.08)	9.4(±2.6)	.90(±.02)	.10(±.05)	20(±3.6)
100	.99(±.005)	.20(±.06)	1.2(±1.1)	.991(±.007)	.03(±.03)	1.8(±1.3)
200	.999(±.0003)	.20(±.05)	0(±0)	.9999(±.0005)	.01(±.01)	.01(±.10)
Band 2	AUC	FPR(%)	FNR (%)	AUC	FPR (%)	FNR (%)
$n = 50$	.82(±.01)	.10(±.05)	36.7(±1.5)	.803(±.008)	.02(±.02)	39.3(±1.5)
100	.84(±.01)	.11(±.04)	32.7(±1.9)	.83(±.01)	.03(±.02)	34.9(±1.6)
200	.93(±.01)	.11(±.04)	14.2(±2.4)	.82(±.01)	.03(±.02)	36.7(±1.8)

Table 5: Performance measures for L_0 regularized regression for graphical structure detection over 100 simulations, where values in the parenthesis are the standard deviations.

BIC is a heavier penalty and forces most of the weaker correlations with $a_{ij} = 0.25$ to 0. In addition, BIC has slightly larger AUCs for Band 1 network with strong correlation $r = 0.6$ and larger sample size (n=100, 200). One interesting observation is that the FDRs of both AIC and BIC are well controlled. The maximal FDRs of AIC for the Band 1 and 2 networks are 0.29% and 0.2%, while the maximal FDRs of BIC are only 0.1%, and 0.03% respectively. Controlling false discovery rates is crucial for identifying true associations with high-dimensional data in bioinformatics. In general, AUC increases and both FDR and FNR decrease, as the sample sizes become larger, except for Band 2 network with BIC. The performance of BIC is not necessary better with large sample size, since the penalty λ increases with the sample size.

4 Real Application

The purpose of this application is to identify subnetworks and study the biological mechanisms of potential prognostic biomarkers for ovarian cancer with multi-source gene expression data. The ovarian cancer data was downloaded from the KMplot website(www.kmplot.com/ovar) (Gyorffy et al.2012). They originally got the data from searching Gene Expression Omnibus (GEO; http://www.ncbi.nlm.nih.gov/geo/) and The Cancer Genome Atlas (TCGA;

http:// cancergenome.nih.gov) with multiple platforms. All collected datasets have raw gene expression data, survival information, and at least 20 patients available. They merged the datasets across different platforms carefully. The final data has 1287 patients samples, and 22277 probe sets representing 13435 common genes. We identified 112 top genes that are associated with patient survival times using univariate COX Regression. We constructed a co-expression network from the 112 genes with L_0 regularized regression and identified biologically meaningful subnetworks (modules) associated with patient survival. Network is constructed with positive correlation only and BIC. The computational time for constructing such network is less than 2 seconds. One survival associated subnetwork we identified is given in Figure 2. The 22 genes on the subnetwork were then uploaded onto STRING

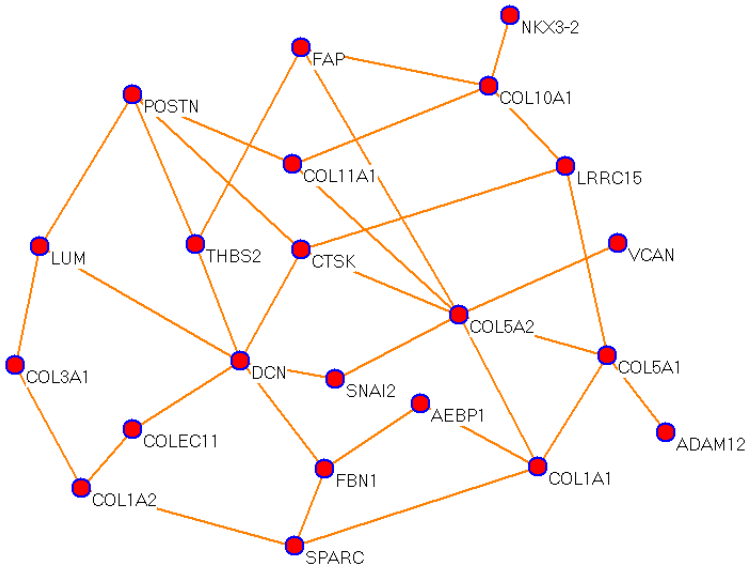

Figure 2: Subnetwork constructed with L_0 penalized regression, multi-source gene expression profiling, and BIC

(http://string-db.org/). STRING is an online database for exploring known and predicted protein-protein interactions (PPI). The interactions include direct (physical) and indirect (functional) associations. The predicted methods for PPI implemented in STRING include text mining, national databases, experiments, co-expression, co-occurrence, gene fusion, and neighborhood on the chromosome. The PPI network for the 22 genes are presented in Figure 3. Comparing Figure 3 and Figure 2, We conclude that the 22 identified genes on the subnetwork of Figure 2 are functioning together and have enriched important biological in-

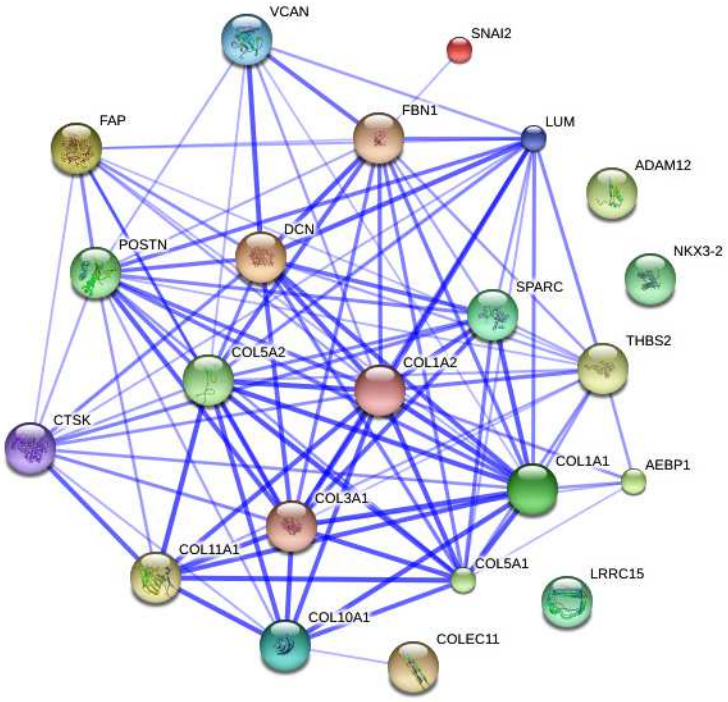

Figure 3: Known and predicted protein protein interactions with the 22 genes on the subnetwork of Figure 2, where nodes represent proteins (genes) and edges indicate the direct (physical) and indirect (functional) associations. Stronger associations are represented by thicker lines.

teractions and associations. Ninteen out of 22 genes on the survival associated subnetwork also have interactions on the known and predicted PPI network, except for genes LRRC15, ADAM12, and NKX3-2. Even though they are not completely identical, many interactions on our subnetwork can also be verified on the PPI interaction network of Figure 3. For instance, collagen COL5A2 is the most important genes with the largest number of degrees (7) on our subnetwork. Six out of 7 genes that link to COL5A2 also have direct edges on the PPI network. Those direct connected genes (proteins) include FAP, CTSK, VCAN, COL1A1, COL5A1, and COL11A1. The remaining gene SNAI2 was indirectly linked to COL5A2 through FBN1 on the PPI network. In addition, one of the other important genes with the degree of the node (6) is Decorin (DCN). 4 out of 6 genes directly connected to DCN on our subnetwork were confirmed on the PPI network, including FBN1, CTSK, LUM, and THBS2. The remain two genes (SNAI2, and COLEC11) are indirectly connected to DCN on the PPI network. As indicated on Figure 2, the remaining 5 important genes with degree of node 4 are POSTN, CTSK, COL1A1, COL5A1, and COL10A1, and 8 genes with

degree of node 3 are FBN1, LUM, LRRC15, COL11A1, THBS2, SPARC, COL1A2, and FAP, respectively. Furthermore, those 22 genes are involved in the biological process of GO terms, including extracellular matrix organization and disassembly and collagen catabolic, fibril, and metabolic processes. They are also involved in several important KEGG pathways including ECM-receptor interaction, Protein Digestion and Absorption, Amoebiasis, Focal Adhesion, and TGF-beta Signaling pathways. Finally, a large proportion of the 22 genes are known to be associated with poor overall survival (OS) in ovarian cancer. For instance, VCAN and POSTN were demonstrated in *vitro* to be involved in ovarian cancer invasion induced by TGF-β signaling (Yeung et al., 2013), and COL11A1 was shown to increase continuously during ovarian cancer progression and to be highly over-expressed in recurrent metastases. Knockdown of COL11A1 reduces migration, invasion, and tumor progression in mice (Cheon et al.2014). Other genes such as FAP, CTSK, FBN1, THBS2, SPARC, and COL1A1 are also known to be ovarian cancer associated (Riester et al., 2014; Zhao et al., 2011; Zhang et al., 2013; Gardi et al., 2014; Tang & Feng 2014; Yu et al., 2014). Those genes contribute to cell migration and the progression of tumors and may be potential therapeutic targets for ovarian cancer. Further studies with the rest of the genes on the subnetwork are required to explore their biological mechanisms and potential clinical applications.

5 Conclusions

We proposed an efficient EM algorithm for variable selection with L_0 regularized regression. The proposed algorithm finds the optimal solutions of L_0, through solving a sequence of L_2 based ridge regressions. Given an initial solution, the algorithm will be guaranteed to converge to a unique solution under mild conditions, and the EM algorithm will be closer to the optimal solution after each iteration. Asymptotic properties, namely consistency and oracle properties are established under mild conditions. Our method apply to fixed, diverging, and ultra-high dimensional problems. We compare the performance of L_0 regularized regression and LASSO with simulated low and high dimensional data. L_0 regularized regression outperforms LASSO by a substantial margin under different correlation structures. Unlike LASSO, which selects more features than necessary, L_0 regularized regression chooses the true model

with high accuracy, less bias, and smaller test MSE, especially when the correlation is weak. Cross-validation with the computation of the entire regularization path is computationally intensive and time consuming. Fortunately L_0 regularized regression does not require it. The optimal λ_{opt} can be directly determined from AIC, BIC, and RIC. Those criteria are optimal under appropriate conditions. We demonstrate that both AIC and BIC performed well when compared to cross-validation. Therefore, there is a big computational advantage of L_0, especially with BIGDATA. In addition, We demonstrate that L_0 regularized regression controls the false discovery (positive) rate (FDR) well with both AIC and BIC with the simulation of graphical models. The FDR is very low under different sample sizes with both AIC and BIC. Controlling FDR is crucial for biomarker discovery and computational biology, because further verifying the candidate biomarkers is time-consuming and costly. We applied our proposed method to construct a network for ovarian cancer from multi-source gene-expression data, and identified a subnetwork that is important both biologically and clinically. We demonstrated that we can identify biologically important genes and pathways efficiently. Even though we demonstrated our method with gene expression data, the proposed method can be used for RNA-seq, and metagenomic data, given that the data are appropriately normalized.

Appendix

The proposed approach for L_0 regularized regression method can be extended to a general L_p $p \in [0, 2]$ naturally. Mathematically, the general L_p problem can be defined as:

$$E = \frac{1}{2}||\mathbf{y} - X\theta||^2 + \frac{\lambda}{2}\sum_{j=1}^{m}|\theta|^p,$$

which is equivalent to

$$E = \frac{1}{2}||\mathbf{y} - X\theta||^2 + \frac{\lambda}{2}\sum_{j\in m}\frac{\theta_j^2}{\eta_j^{2-p}}$$

$$\eta = \theta.$$

Similar ideas in the manuscript can be used to get the the following equation for the general L_pEM method:

$$\eta^{2-p} \odot \frac{\partial E}{\partial \theta} = \lambda\theta - \eta^{2-p} \odot X^t(\mathbf{y} - X\theta) = \lambda\theta - X_\eta^t(\mathbf{y} - X\theta) = 0,$$

where $X_\eta^t = [\eta^{2-p} \odot \mathbf{x}_1^t, \ldots, \eta^{2-p} \odot \mathbf{x}_n^t]$. Solving Equation (9), we have the following explicit solution.

$$\theta = (X_\eta^t X + \lambda I)^{-1} X_\eta^t \mathbf{y}$$
$$\eta = \theta,$$

The general L_pEM algorithm is as follows:

L_pEM Algorithm:

Given a $0 < \lambda \leq \lambda_{\max}$,and $p \in [0, 2]$, small numbers ϵ and ε,
and training data $\{X, \mathbf{y}\}$,
Initializing $\theta = (X^t X + \lambda I)^{-1} X^t \mathbf{y}$,
While 1,
  E-step: $\eta = \theta$
  M-step: $X_\eta^t = \eta^{2-p} \odot X^t = [\eta^{2-p} \odot \mathbf{x}_1, \ldots, \eta^{2-p} \odot \mathbf{x}_n^t]$
    $\theta = (X_\eta^t X + \lambda I)^{-1} X_\eta^t \mathbf{y}$
  if $||\theta - \eta|| < \varepsilon$, Break; End
End
$\theta(|\theta| < \epsilon) = 0$.

References

Akaike, H. (1974). A new look at the statistical model identification. IEEE T. Automat. Contr. 19, 716723.

Bunea, F., Tsybakov, A, and Wegkamp, M. (2007). Sparsity oracle inequalities for the Lasso. *Electron. J. Stat.*, **1**, 169-194.

Candes, E.J. and Plan, Y. (2009). Near-ideal model selection by l_1 minimization. *Ann. Statist.*, **37**(5A), 2145-2177.

Cheon DJ, Tong Y, Sim MS, Dering J, Berel D, Cui X, Lester J, Beach JA, Tighiouart M, Walts AE, Karlan BY, Orsulic S. (2014), A collagen-remodeling gene signature regulated by TGF- signaling is associated with metastasis and poor survival in serous ovarian cancer. Clin Cancer Res. 2014 Feb 1;20(3):711-23. doi: 10.1158/1078-0432.

Chi Z. (2009). L_0 regularized estimation for nonlinear models that have sparse underlying linear structures. arXiv:0910.2517v1 [math.ST] 14 Oct 2009.

Dicker, L., Huang, B. and Lin, X. (2012). Variable selection and estimation with the seamless-L_0 penalty. Statistica Sinica. In press. doi: 10.5705/ss.2011.074.

Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. J. Am. Stat. Assoc. 96, 13481361.

Foster, D. and George, E. (1994). The risk inflation criterion for multiple regression. Ann. Statist. 22, 19471975.

Gardi NL, Deshpande TU, Kamble SC, Budhe SR, Bapat SA. (2014), Discrete molecular classes of ovarian cancer suggestive of unique mechanisms of transformation and metastases. Clin Cancer Res. 2014 Jan 1;20(1):87-99.

Gyorffy B, Lánczky A, Szállási Z. (2012), Implementing an online tool for genome-wide validation of survival-associated biomarkers in ovarian-cancer using microarray data from 1287 patients. Endocr Relat Cancer. 2012 Apr 10;19(2):197-208. doi: 10.1530/ERC-11-0329.

Lee HK, Hsu AK, Sajdak J, Qin J, Pavlidis P (2004), Coexpression analysis of human genes across many microarray data sets. Genome Res. 2004 Jun;14(6):1085-94.

Lin, D., Foster, D. P., & Ungar, L. H. (2010). A risk ratio comparison of l0 and l1 penalized regressions. University of Pennsylvania, Tech. Rep.

Liu H, Roeder K, and Wasserman L. (2010) Stability approach to regularization selection for high dimensional graphical models. Advances in Neural Information Processing Systems, 2010.

Liu Q, Ihler A (2011), Learning scale free networks by reweighted L_1 regularization. AISTATS (2011).

Liu Y, Wu Y (2007), Variable Selection via A Combination of the L_0 and L_1 Penalties, Journal of Computational and Graphical Statistics, 16 (4): 782798.

Liu Z, Lin S, Tan M. (2010) Sparse support vector machines with Lp penalty for biomarker identification. IEEE/ACM Trans Comput Biol Bioinform. 7(1): 100-7.

Lu H and Zhang Y (2013), Sparse Approximation via Penalty Decomposition Methods, SIAM Journal on Optimization, 23(4):2448-2478.

Mancera L and Portilla J. L_0 norm based Sparse Representation through Alternative Projections. In Proc. ICIP, 2006

Mazumder R, Friedman, JH, Hastie T (2011), SparseNet : Coordinate Descent with Non-Convex Penalties, JASA, 2011.

Meinshausen, N. & Bhlmann, P. Stability selection. J. R. Stat. Soc. Series B Stat. Methodol. 72, 417473 (2010).

Peng,J, Wang P, Zhou N, and Zhu J. (2009), Partial correlation estimation by joint sparse regression models. JASA, 104 (486):735-746, 2009.

Riester M, Wei W, Waldron L, Culhane AC, Trippa L, Oliva E, Kim SH, Michor F, Huttenhower C, Parmigiani G, Birrer MJ. (2014), Risk prediction for late-stage ovarian cancer by meta-analysis of 1525 patient samples. J Natl Cancer Inst. 2014 Apr 3;106(5). pii: dju048. doi: 10.1093/jnci/dju048.

Schwarz, G. (1978). Estimating the dimension of a model. Ann. Statist. 6, 461464.

Tang L, Feng J. (2014), SPARC in Tumor Pathophysiology and as a Potential Therapeutic Target. Curr Pharm Des. 2014 Jun 19. [Epub ahead of print].

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. J. Roy. Stat. Soc. B. 58, 267288.

Yeung TL, Leung CS, Wong KK, Samimi G, Thompson MS, Liu J, et al. TGF-beta modulates ovarian cancer invasion by upregulating CAF-derived versican in the tumor microenvironment. Cancer Res 2013;73:501628.

Yu PN, Yan MD, Lai HC, Huang RL, Chou YC, Lin WC, Yeh LT, Lin YW. (2014), Downregulation of miR-29 contributes to cisplatin resistance of ovarian cancer cells. Int J Cancer. 2014 Feb 1;134(3):542-51.

Zhang W, Ota T, Shridhar V, Chien J, Wu B, Kuang R. (2013), Network-based survival analysis reveals subnetwork signatures for predicting outcomes of ovarian cancer treatment. PLoS Comput Biol. 2013;9(3):e1002975. doi: 10.1371/journal.pcbi.1002975.

Zhao G, Chen J, Deng Y, Gao F, Zhu J, Feng Z, Lv X, Zhao Z (2011), Identification of NDRG1-regulated genes associated with invasive potential in cervical and ovarian cancer cells. Biochem Biophys Res Commun. 2011 Apr 29;408(1):154-9.

Zou, H. (2006). The adaptive lasso and its oracle properties. J. Am. Stat. Assoc. 101, 14181429.

Zou, H. and Zhang, H. (2009). On the adaptive elastic-net with a diverging number of parameters. Ann. Statist. 37, 17331751.

Zhang, C. (2010). Nearly unbiased variable selection under minimax concave penalty. Ann. Statist. 38, 894942.